CLEAN VERSION OF CHANGES TO THE SPECIFICATION, ABSTRACT, AND ALL PENDING CLAIMS

In the Specification:

(1) Please replace the fourth full paragraph at page 4, lines 23-28 with the following:

A1

The inference engine can be employed to infer informational goals from the user input. The inputs to the inference engine can include, but are not limited to, the user input (parsed and/or unparsed), extrinsic data, and information retrieved from the inference model. The answer generator may be employed to produce an answer to a query. The inputs to the answer generator may include, but are not limited to, the original user input, extrinsic data and informational goals inferred by the inference engine.

(2) Please replace the paragraph beginning at page 7, line 20, and ending at page 8, line 2, with the following:

A2

Referring initially to Fig. 1, a schematic block diagram illustrates a system 100 for inferring informational goals in queries, which may be used for enhancing responses to queries presented to an information retrieval system (*e.g.*, a question answering system). The system 100 includes a query subsystem 110 that is employed in processing a user input 120 and extrinsic data 130 to produce an output 180. The user input 120 can be, for example, a query presented to a question answering application. The extrinsic data 130 can include, but is not limited to, user data (*e.g.*, applications employed to produce query, device employed to generate query, current content being displayed), context (*e.g.*, time of day, location from which query was generated, original language of query) and prior query interaction behavior (*e.g.*, use of query by example (QBE), use of query/result feedback). The output 180 may include, but is not limited to, one or more responses, an answer responsive to a query in the user input 120, one or more re-phrased queries, one or more suggested queries (that may be employed, for example, in a QBE system) and/or an error code.

(3) Please replace the first full paragraph at page 8, lines 3-14, with the following:

A3

In an exemplary aspect of the present invention, when the output 180 takes the form of one or more responses, the one or more responses may be further processed to vary in length, precision and detail based, at least in part, on the inferred informational goals associated with the query that produced the one or more responses. In another exemplary aspect of the present invention, the output 180 may be subjected to further processing. For example, if the output 180 takes the form of two or more responses, then the responses may be ranked by a ranking process to indicate, for example, the predicted relevance of the two or more responses. Similarly, the output 180 may be further processed by a text focusing process that may examine the output 180 to facilitate locating and displaying the piece(s) of information most relevant to the query. Further, the output 180 may be processed, for example, by a diagramming process that displays information graphically, rather than textually.

(4) Please replace the third full paragraph at page 8, lines 21-31, with the following:

A4

The query subsystem 110 can include an inference engine 112 and a response generator 114. The query subsystem 110 can also receive the user input 120 *via* a natural language processor 116. The natural language processor 116 can be employed to parse queries in the user input 120 into parts that can be employed in predicting informational goals. The parts may be referred to as "observable linguistic features". By way of illustration, the natural language processor 116 can parse a query into parts of speech (*e.g.*, adjectival phrases, adverbial phrases, noun phrases, verb phrases, prepositional phrases) and logical forms. Structural features including, but not limited to, the number of distinct parts of speech in a query, whether the main noun in a query is singular/plural, which noun (if any) is a proper noun and the part of speech of the head verb post modifier can also be extracted from output produced by the natural language processor 116.

(5) Please replace the second full paragraph at page ten, lines 18-26, with the following:

The query subsystem 110 can also include a response generator 114. The response generator 114 can, for example, receive as input predictions concerning informational goals and can access, for example, the knowledge data store 170 to retrieve information responsive to the query in the user input 120. The response generator may also produce responses that are not answers, but that include rephrased queries and/or suggested queries. For example, the query subsystem 110 may determine that the amount and/or type of information sought in a query is so broad and/or voluminous that refining the query is appropriate. Thus, the response generator 114 may provide suggestions for refining the query as the response to the query rather than producing an answer.

A5

(6) Please replace the paragraph beginning at page 11, line 29, and ending at page 12, line 29, with the following:

The learning system 150 can employ both automated and manual means for performing supervised learning, with the supervised learning being employed to construct and/or adapt data structures including, but not limited to, decision trees in the inference model 160. Such data structures can subsequently be employed by the inference engine 112 to predict informational goals in a query in the user input 120. Predicting the informational goals may enhance the response to a query by returning a precise answer and/or related information rather than returning a document as is commonly practiced in conventional information retrieval systems. By way of illustration, the present invention may provide answers of varying length and level of detail as appropriate to a query. In this manner, an exemplary aspect of the present invention may model the expertise of a skilled reference librarian who can not only provide the requested answer but understand the subtleties and nuances in a question, and identify an "appropriate" answer to provide to the querying user. For example, presented with the query "What is the capital of Poland?" traditional question answering systems may seek to locate documents containing the terms "capital" and "Poland" and then return one or more documents that contain the terms "capital" and "Poland". The information consumer may then be forced to read the one or more

A6

documents containing the terms to determine if the answer was retrieved, and if so, what the answer is. The present invention, by inferring informational goals, identifies conditions under which a more extended reply, such as "Warsaw is the capital and largest city of Poland, with a population of approximately 1,700,00" is returned to the user. The present invention may, for example, set values for several variables employed in analyzing the query (*e.g.,* Information Need set to "Attribute"; Topic set to "Poland"; Focus set to "capital"; Cover Wanted set to "Precise", and Cover Would Give set to "Additional"). Further, the present invention may determine that pictures of landmarks, a city street map, weather information and flight information to and from Warsaw may be included in an appropriate reply. These informational goals are predicted by analyzing the observable linguistic features found in the query and retrieving conditional probabilities that certain informational goals exist from the inference model 160 based on those observable linguistic features. The inference model 160 can be constructed by employing supervised learning with statistical analysis on queries found in one or more query logs 140. The inference model 160 can then be employed by a "run time system" to facilitate such enhanced responses.

A6

(7) Please replace the paragraph beginning at page 12, line 30 and ending at page 13, line 13, with the following:

In one example of the present invention, the learning system 150 and/or the inference engine 112 may further be adapted to control and/or guide a dialog that can be employed to clarify information associated with informational goals, desired level of detail, age and so on. By way of illustration and not limitation, the learning system 150 may make an inference (*e.g.*, age), but then may present a user interface dialog that to facilitates clarifying the age of the user. Thus, the learning system 150 may be adapted, *in-situ*, to acquire more accurate information concerning inferences, with resulting increases in accuracy. Such increased accuracy may be important, for example, in complying with Federal Regulations (*e.g.*, Children's Online Privacy Protection Act). By way of further illustration, the inference engine 112 may make an inference (*e.g.*, level of detail in answer), but then may present a user interface that facilitates clarifying the desired level of detail in an answer. Thus, the inference engine 112 may adapt processes employed in

A7

A7

generating an inference, and may further adapt search and retrieval processes and/or post-search filtering processes to provide a closer fit between returned information and desired coverage.

(8) Please replace the paragraph beginning at page 13, line 21, and ending at page 14, line 6, with the following:

A8

The run time system 200 receives data from a user input 220 and may also receive an extrinsic data 230. The user input 220 can include one or more queries for information. The run time system 200 may receive queries directly and/or may receive parse data from a natural language processor 216. The queries may appear simple (*e.g.*, what is the deepest lake in Canada?) but may contain informational goals that can be employed to enhance the response to the query. For example, the query "what is the deepest lake in Canada?" may indicate that the user could benefit from receiving a list of the ten deepest lakes in Canada, the ten shallowest lakes in Canada, the ten deepest lakes in neighboring countries, the ten deepest lakes in the world and the ten deepest spots in the ocean. While there are time and processing costs associated with inferring the informational goals, retrieving the information and presenting the information to the information consumer, the benefit of providing information rather than documents can outweigh that cost, producing an enhanced information gathering experience.

(9) Please replace the first full paragraph on page fourteen, lines 7-20, with the following:

A9

To facilitate enhancing the informational retrieval experience, the run time system 200 may also examine extrinsic data 230. The extrinsic data 230 can include, but is not limited to, user data (*e.g.*, applications employed to produce query, device employed to generate query, current content being displayed), context (*e.g.*, time of day, location from which query was generated, original language of query) and prior query interaction behavior (*e.g.*, use of query by example (QBE), use of query/result feedback). The user data (*e.g.*, device generating query) can provide information that may be employed in determining what type and how much information should be retrieved. By way of illustration, if the device generating the query is a personal computer, then a first type and amount of information may be retrieved and presented, but if the

A9

device generating the query is a cellular telephone, then a second type and amount of information may be retrieved and presented. Thus, the informational goals of the user may be inferred not only from the observable linguistic features of a query, but also from extrinsic data 230 associated with the query.

(10) Please replace the paragraph beginning at page 14, line 21, and ending at page 15, line 10, with the following:

A10

The run time system 200 includes a query subsystem 210, which in turn includes an inference engine 212 and a response generator 214. The query subsystem 210 accepts parse data produced by a natural language processor 216. The natural language processor 216 takes an input query and produces parse data including, but not limited to, one or more parse trees, information concerning the nature of and relationships between linguistic components in the query (*e.g.*, adjectival phrases, adverbial phrases, noun phrases, verb phrases, prepositional phrases), and logical forms. The query subsystem 210 subsequently extracts structural features (*e.g.*, number of distinct points of speech in a query, whether the main noun in a query is singular/plural, which noun (if any) is a proper noun and the part of speech of the head verb post modifier) from the output of the natural language processor 216. Such parse data can then be employed by the inference engine 212 to, for example, determine which, if any, of one or more data structures in the inference model 240 to access. By way of illustration, first parse data indicating that a first number of nouns are present in a first query may lead the inference engine 212 to access a first data structure in the inference model 240 while second parse data indicating that a certain head verb post modifier is present in a second query may lead the inference engine 212 to access a second data structure in the inference model 240. By way of further illustration, the number of nouns and the head verb post-modifier may guide the initial access to different decision trees (*e.g.*, one for determining information need, one for determining focus) and/or the number of nouns and the head verb post-modifier may guide the access to successive sub-trees of the same decision tree.

(11) Please replace the paragraph beginning on page 19, line 26, and ending at page 20, line 6, with the following:

Based on the high level information goals 530 inferred from the observable linguistic features in the query 520, one or more content sources (*e.g.*, content source 550_{A1}, 550_{A2} through 550_{An}, n being an integer, referred to collectively as the content sources 550) may be accessed by the query system 540 to produce a response 560 to the query 520. The content sources 550 can be, for example, online information sources (*e.g.*, newspapers, legal information sources, CD based encyclopedias). Based on the informational goals 530 inferred from the query 520, the query system 540 can return information retrieved from the content sources 550. Further, the information may vary in aspects including, but not limited to, content, length, scope and abstraction level, for example, again providing an improvement over conventional systems.

A11

(12) Please replace the first full paragraph at page 20, lines 7-25, with the following:

Referring now to Fig. 6 a training system 600 including a natural language processor 640, a supervised learning system 660 and a Bayesian statistical analyzer 662 is illustrated. The training system 600 includes a question store 610 as a source of one or more sets of questions suitable for posing to a question answering system. The question store 610 may be a data store and/or a manual store. The question store 610 may be configured to facilitate specific learning goals (*e.g.*, localization). By way of illustration, questions posed from a certain location (*e.g.*, Ontario) during a period of time (*e.g.*, Grey Cup Week) to an online question answering service may be stored in the question store 610. The questions may be examined by a question examiner (*e.g.*, linguist, cognitive scientist, statistician, mathematician, computer scientist) to determine question suitability for training, with some questions being discarded. Further, the questions in the question store 610 may be selectively partitioned into subsets including a training data subset 620 and a test data subset 630. In one example aspect of the present invention, questions in the question store 610, the training data 620 and/or the test data 630 may be annotated with additional information. For example, a linguist may observe linguistic features and annotate a question with such human observed linguistic features to facilitate evaluating the operation of the

A12